

(12) UK Patent Application (19) GB (11) 2 326 255 (13) A

(43) Date of A Publication 16.12.1998

(21) Application No 9806749.7

(22) Date of Filing 25.04.1998

(30) Priority Data

(31) 08852263 (32) 07.05.1997 (33) US

(71) Applicant(s)

International Business Machines Corporation
(Incorporated in USA - New York)
Armonk, New York 10504, United States of America

(72) Inventor(s)

David Louis Kaminsky
Richard Adam King
Zhiyong Li

(74) Agent and/or Address for Service

G M Zerbi
IBM United Kingdom Limited, Intellectual Property
Department, Mail Point 110, Hursley Park,
WINCHESTER, Hampshire, SO21 2JN,
United Kingdom

(51) INT CL⁶

G06F 9/46

(52) UK CL (Edition P)

G4A APX

(56) Documents Cited

GB 2305270 A US 5684955 A

(58) Field of Search

UK CL (Edition P) G4A APX

INT CL⁶ G06F 9/46

Online: WPI, INSPEC, COMPUTER

(54) Abstract Title

Automatic object distribution enables remote running of objects using local method calls

(57) Object oriented programs are run as distributed programs without any explicit networking code, and without using an interface definition language (IDL). When objects Y 303 and X 305 are distributed between memories 301 and 302, two proxies, X' 315 and X'' 317, are dynamically generated, containing classes serving as proxies for X. When object Y makes a call, 320, against a method in the original class X, that call is actually made against a method in the new class X residing locally and contained in file X'. X' passes the call information, 321, to X'', which makes a local call, 322, to X. The results are propagated back to Y by way of X'' and X'. Thus the proxies allow calls to flow across a network as if they were local, and programmers can experiment with different object distributions without complicating the programming task.









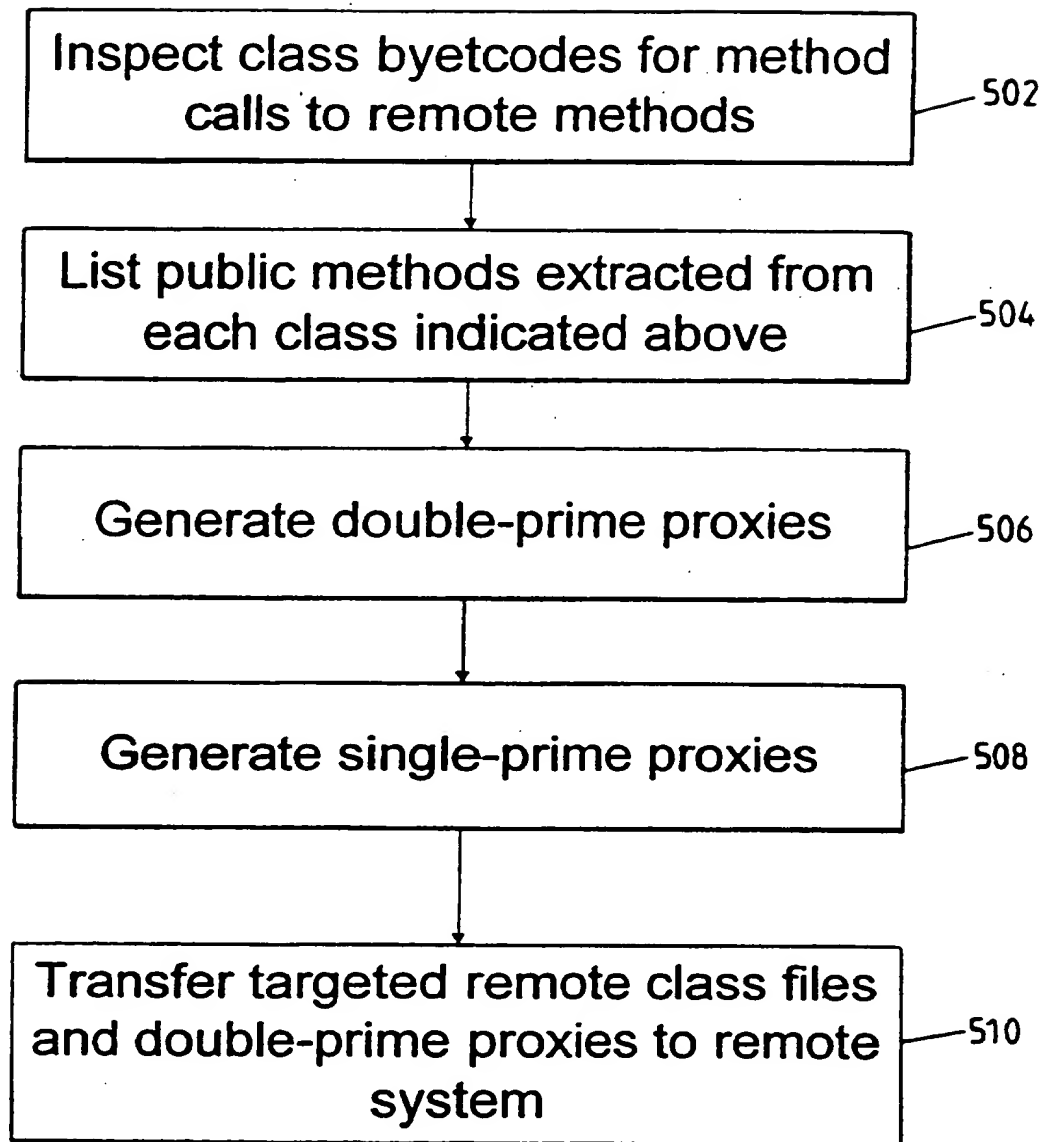
FIG. 4B

GB 2 326 255 A

FIG. 1

FIG. 2FIG. 3

FIG. 4AFIG. 4B

FIG. 5

A PROCESS FOR RUNNING OBJECTS REMOTELY

Field of the Invention

5 The present invention relates generally to computer programming methods and systems, and, in particular, to object oriented programming and to methods and systems of running object oriented programs on multiple computers connected by a network.

Background of the Invention

10 The distribution of a single program, so that a portion of the program executes on more than one computer has become more pervasive as "desktop" computers become more powerful. Figure 1 demonstrates an
15 example of a computer network containing several distributed computers upon which an application can execute. While most computer networks are many orders of magnitude larger, this small network is used as an example. Presently, many computer systems allow objects to communicate over a network. One method of allowing objects to communicate over a
20 network is RPC, which stands for Remote Procedure Calls. RPC has existed since the mid 1980s and is further described by Birrell and Nelson in Implementing Remote Procedure Calls, ACM Transactions on Computer Systems 2, 1984, pp 39-59. Further improvements have been made to RPC, such as DSOM, or Distributed Systems Object Model, continuing into the present.
25 In the systems of the prior art, the programmer describes the interfaces to the objects using an Interface Definition Language (IDL). The systems then provide tools that automatically transform the IDL specification into executable code.

30 Recently, Sun Microsystems released an object oriented programming language called Java which includes a capability similar to that of RPC and DSOM called Remote Method Invocation, or RMI. Using systems which are Java enabled, the programmer can now write a distributed object
35 program without explicit recognition of the network upon which the program will be running. However, using this technology, the programmer or implementer is still required to write additional code that enables remote invocation of objects (or write the IDL specification). As with all coding tasks, writing IDL is error-prone. In addition, should the
40 programmer determine that the point at which he chose to distribute the objects is sub-optimal, he must delete the current specification, and write another. Since choosing the optimal distribution point is a

difficult task, this problem occurs often in modern programming environments. In addition, the optimal distribution point may vary based on computing topologies. For example, if two computers are equally powerful, then the optimal distribution might have an equal number of objects executing on each computer; if the two computers have disparate capabilities, then the more powerful computer might execute more of the objects. Since heterogeneous computing topologies are common, code written statically using an IDL often results in suboptimal performance.

Alternatively, many program-creation tools allow a programmer to specify a program using the tool, and indicate the object boundaries at which the objects are to be distributed. The tool then generates the executable code for the program. However, these tools generate static program partitions, and thus suffer from heterogeneity problems similar to IDL systems. In addition, tools do not typically interoperate. Thus, multiple programmers working on the same program are forced to standardize on the same tool to work effectively together. Since tools have strengths and weaknesses, standardizing on a single tool is undesirable.

Summary of the Invention

It is an object of the present invention to provide a technique which overcomes the above drawbacks.

According to the present invention we provide a method for distributing objects of a program, each object containing one or more programmed methods, across more than one physical device, said method comprising the steps of: identifying all of the objects in the program; determining which of the objects are to reside on a first computer and which of the objects are to reside on a second computer; identifying all programmed methods contained in each object that may be accessed from a remote computer; generating a first proxy and a second proxy for each object that may be accessed from a remote computer, said first proxy residing on said first computer and said second proxy residing on said second computer, said first proxy containing network linkage and indication to access said programmed methods on said second computer and said second proxy containing network linkage and indication to access said programmed methods on said first computer; and, accessing said remote programmed methods through said proxies.

Further according to the present invention we provide a computer-readable storage medium having computer-readable program code means embodied in said medium for distributing objects of a program, each object containing one or more programmed methods, across more than one physical device, said computer-readable program code means comprising:

5 computer-readable program code means for identifying all of the objects in the program; computer-readable program code means for determining which of the objects are to reside on a first computer and which of the objects are to reside on a second computer; computer-readable program

10 code means for identifying all programmed methods contained in each object that may be accessed from a remote computer; computer-readable program code means for generating a first proxy and a second proxy for each object that may be accessed from a remote computer, said first proxy residing on said first computer and said second proxy residing on said

15 second computer, said first proxy containing network linkage and indication to access said programmed methods on said second computer and said second proxy containing network linkage and indication to access said programmed methods on said first computer; and, computer-readable program code means for accessing said remote programmed methods through

20 said proxies.

Also according to the present invention we provide a system for distributing objects of a program, each object containing one or more programmed methods, across more than one physical device, said system comprising:

25 means for identifying all of the objects in the program; means for determining which of the objects are to reside on a first computer and which of the objects are to reside on a second computer; means for identifying all programmed methods contained in each object that may be accessed from a remote computer; means for generating a

30 first proxy and a second proxy for each object that may be accessed from a remote computer, said first proxy residing on said first computer and said second proxy residing on said second computer, said first proxy containing network linkage and indication to access said programmed methods on said second computer and said second proxy containing network

35 linkage and indication to access said programmed methods on said first computer; and, means for accessing said remote programmed methods through said proxies.

40 A preferred embodiment of the present invention describes a system, method and program product for executing or running objects remotely which is herein referred to as Automatic Object Distribution (AOD). As

shown in Figure 2, an object oriented program is written with many independent objects that call each other to perform a unified function. Figure 2 depicts a computer memory 101 in which objects 102 V, W, X, Y & Z all reside. The system according to a preferred embodiment of the present invention allows a user to move some of the objects from a first computer to a second computer to be executed. This is shown in Figure 3. In Figure 3 objects V & Y 205 remain in a first computer memory 201 while objects W, X, & Z 207 are moved to a second computer memory 203. This can be done, for instance, to balance the workload between the first computer and the second computer. This creates additional complexities since each of the program objects are capable of calling external routines herein referred to as methods. The separation of the program objects into multiple computers for execution requires that all of the methods or objects which, due to the separation of the objects to different computer memories now access methods which are no longer coresident with the calling method on the same system, know that some of the objects are located elsewhere on the network and understand how to access those remote objects.

The Automatic Object Distribution (AOD) according to a preferred embodiment of the present invention eliminates the requirement that programmers write their programs with the knowledge that the objects will be distributed over a network. It also eliminates the need for programmers to write programs in a special language that supports distributed objects, or to write using any particular tool when wanting to distribute the work across the network. Instead, programmers write programs in a common object-oriented language, such as Java (Java is a trademark of Sun Microsystems), and compile the code into unlinked executable code. (In the computing literature, unlinked executable code is often called "object code". When describing object-oriented systems, the use of the term "object code" creates confusion, therefore the term "bytecodes" has been adopted for purposes of the present invention.) The programmer writes the code exactly as if writing a program that will execute on a single machine; he does not decorate the program with any additional information as he would with IDL systems. The programmer then decides where the distribution should occur, and the AOD generates the distribution code.

The AOD performs the distribution by generating two proxies. The proxies allow method calls written for local invocation to be invoked over a network. If a split is indicated between some class Y to be

executed locally, and some class X to be executed remotely, the AOD process generates the proxies to overcome the intervening network. Together, these proxies intercept the calls from Y, pass them to X, and return the result to Y. This is shown in more detail in Figures 4A and 4B.

Figure 4A represents the computer system in which a program is running in a single memory 301. Both object Y 303 and object X 305 reside on the same memory and a method call 311 is used to invoke object X 305 from object Y 303.

As shown in Figure 4B, the AOD process uses information in the bytecode file for a class, say X 305, to generate the two proxy files (X' 315 and X" 317) that contain classes serving as proxies for X 305. X' 315 contains a class named X that has all of the public methods in X 305. (Note that since the original class X 305 file will be located on the remote machine, no name conflict exists.) X" 317 contains a class with a unique name (X") that contains the method that makes calls to the public methods in class X. X' 315 resides on the same machine as Y 303; X" 317 resides on the same machine as X 305. Since X' 315 contains a class called X and that class contains all of the public methods in X 305, when Y 303 makes a call against a method in the original class X 305, that call is actually made against a method in the new class X residing locally and contained in file X' 315. X' 315 passes the call information to X" 317 which then makes a local call to X 305. The results of the local call are propagated back to Y 303 by way of X" 317 and X' 315. The result is that the application, written to execute locally, can be distributed across the network without requiring that the programmer write or change any of the program code.

Brief Description of the Drawings

Figure 1 depicts a sample network upon which the present invention may appear.

Figure 2 depicts a computer memory containing five objects.

Figure 3 depicts a distributed system upon which two objects execute in one computer memory and three objects execute in a different computer memory.

Figure 4A depicts a method where two objects execute in the same memory.

Figure 4B depicts the objects of Figure 4A after being distributed using the present invention.

Figure 5 is a flow chart of the distribution method of the present invention.

Detailed Description of the Preferred Embodiment

The preferred embodiment of the present invention will be herein described in more detail with reference to the drawings. The present invention is not limited to the preferred embodiment. The present invention is applicable to any automated distributed programming utilizing proxies. Automatic object distribution (AOD) of the preferred embodiment uses the Java programming environment although the present invention is not limited to Java and the application of this invention to other object-oriented environments would be straightforward to one skilled in the art. The preferred embodiment of the present invention assumes that the programmer has written a program in Java as if it were to run locally (e.g., it does not include any Remote Method Invocation code), and used the standard Java compiler to compile the source files into corresponding bytecode files. The AOD process then utilizes the convention that, in object oriented programming, all objects are to be accessed through method calls (not through direct access to public variables). In Java, each class compiles into a separate bytecode file. (Hence, bytecode files are often called "class files.") The programmer then identifies which classes are to be moved to a remote machine and which will remain local. The process of identifying such classes can be accomplished by typing the names of the files containing the classes into a tool implementing the present invention.

To run the objects remotely using the preferred embodiment of the present invention, the AOD, using the input as to which classes will be kept local and which classes will be moved remote, does the following as referred to in Figure 5:

- 1) AOD inspects the bytecode for each class to determine which classes will receive method calls from classes that will execute remotely 502.

2) For each object X that will be called upon from a remote machine, a list of its public methods is extracted from bytecode file for X, and stored 504.

5 3) Based on the list of public methods, a proxy, called X* is generated. X* is constructed to contain methods which have similar names as the public methods in X. Each of the constructed methods make corresponding method calls against X. It also contains the code necessary to register itself with the RMI registry 506.

10 4) Based on the list of public methods, a proxy, called X' is generated. X' is constructed to contain the public methods in X. Each of the constructed methods make corresponding remote method calls (using Java's RMI function) against X*. This proxy class has the same name that X does 508.

15 5) All of the original bytecode files designated as remote and all "double-prime proxies" are transferred to the remote machine; all bytecode files designated as local and all "prime" proxies remain on the local machine 510.

Each step is described in more detail below:

25 1) The Java compiler (and compilers of many object-oriented languages) generates a bytecode file in a well-known format which allows the inspection process to occur successfully. To allow the various bytecode files to be linked together, each bytecode file contains information about the methods called from each class. Thus, by reading the information about the methods from the file, 30 AOD can generate the required list of interrelated classes.

35 2) Similarly, the bytecode files contain information about the public methods of the class described by the bytecode file. Thus, by reading the information about the public methods from the file, AOD can generate a list of all public methods. The bytecode files also contain information about parameters and return values.

40 Note that some implementations of Java contain a function called "reflection" that allows programmatic determination of public methods supported by a class. If reflection is implemented, it can be used to replace the process of inspecting the bytecode file.

3) Given the list of all public methods on a class, it is straightforward to generate a double-prime proxy class (X*) that contains the similarly named methods, and each of these methods makes calls to the method on X. In the preferred embodiment, the calls are generated as Java source, and compiled dynamically using Java's standard dynamic compilation facility. Thus, if X contains methods A and B, X* would be written as:

```

Class Xprimeprime extends UnicastRemoteObject
                                implements XprimeprimeInterface {
X x;
Xprimeprime () { // initializer
    x = new X ();
}
A__AOD () { // passes through calls to A
    x.A();
}
B__AOD () { // passes through calls to B
    x.B();
}

```

Note that we have used the convention that the method in the proxy is suffixed with "__AOD." For example, the method "A__AOD" calls the method "A" in X. Any convention can be used so long as it is reproducible and consistent.

4) Given the list of all public methods on a class, it is similarly straightforward to generate a prime proxy class (X') that contains those methods, and that makes remote calls (using RMI) to similarly named methods on X*. In the preferred embodiment, the calls are generated as Java RMI source, and compiled dynamically using Java's standard dynamic compilation facility.

In addition to the proxy X*, an interface X*Interface which contains the signatures for all of the public methods in X is also generated as required by Java's standard RMI implementation.

Given the same information on all public methods on a class, it is also straightforward to generate the interface X*Interface.

```

Interface XprimeprimeInterface extends java.rmi.Remote {
    A__AOD();
    B__AOD ();
}

```

Thus, if X contains A and B, X' would be generated as:

```

Class Xprime {
    XprimeprimeInterface x;

    Xprime () { // constructor
        x = (XprimeprimeInterface)Naming.lookup(Xprimeprime); // do the RMI
        lookup
    }

    A() { // call remote A
        x.A__AOD();
    }

    B() { // call remote B
        x.B__AOD();
    }
}

```

In this example, Naming.lookup() will return a remote reference to X by querying the RMI registry. By convention, the return object is cast to the type of X Interface.

Note that while the proxy class is called X, no name collision occurs with the original X. Since the proxy X will execute on the local machine, and the original X was identified as moving to the remote machine, only one X exists per machine.

5) The files are then transferred by any means available on the system where this function is implemented.

Once the proxies have been established as shown in Figure 4B, the following occurs at execution time:

Assume that the programmer chooses to distribute the program between some object Y and some object X, where Y accesses X only through

method calls. When Y calls a public method `foo()` on X, the following occurs. First Y makes the call to `foo()` as usual. (It must since its code is unchanged). Since the proxy `Xprime` is the only X on the local machine and it contains a method `foo()`, the call will be made on the prime proxy for X.

By construction, the method `foo()` in X makes an RMI call to method `foo__AOD()` on the remote double-prime proxy, returning the results. By construction, `foo().__AOD` on the double-prime proxy makes a local call to `foo()` on X. Since X is unchanged, it will execute the original `foo()`, passing back the results to `foo__AOD` in the double-prime proxy, which passes them to the prime proxy, which returns them to Y. Thus, the call is semantically equivalent to a local call.

Because this process does not impact the original files, and does not require any programmer intervention to realize the distribution, it becomes quite easy for a programmer to test different class distribution points. This enables programs to execute more efficiently without significant impact to the programmers.

In addition, since the distribution requires no additional programming, a system administrator can deploy multiple instances of the same program on different computing configurations, each with its own distribution (and thus performance characteristics), depending on the particular characteristics of the segment of the computing environment. This distribution requires no additional programming and does not require any change to the original program.

CLAIMS

1. A method for distributing objects of a program, each object containing one or more programmed methods, across more than one physical device, said method comprising the steps of:

identifying all of the objects in the program;

determining which of the objects are to reside on a first computer and which of the objects are to reside on a second computer;

identifying all programmed methods contained in each object that may be accessed from a remote computer;

generating a first proxy and a second proxy for each object that may be accessed from a remote computer, said first proxy residing on said first computer and said second proxy residing on said second computer, said first proxy containing network linkage and indication to access said programmed methods on said second computer and said second proxy containing network linkage and indication to access said programmed methods on said first computer; and,

accessing said remote programmed methods through said proxies.

2. A method as claimed in Claim 1 wherein said second proxy contains similarly named methods to the those methods in the objects remaining on said first computer and wherein said first proxy contains similarly named methods to those methods in the objects moved to said second computer.

3. A method as claimed in Claim 1 or 2 wherein said network linkage includes statements for remote invocation of the programmed methods.

4. A computer-readable storage medium having computer-readable program code means embodied in said medium for distributing objects of a program, each object containing one or more programmed methods, across more than one physical device, said computer-readable program code means comprising:

computer-readable program code means for identifying all of the objects in the program;

computer-readable program code means for determining which of the objects are to reside on a first computer and which of the objects are to reside on a second computer;

computer-readable program code means for identifying all programmed methods contained in each object that may be accessed from a remote computer;

computer-readable program code means for generating a first proxy and a second proxy for each object that may be accessed from a remote computer, said first proxy residing on said first computer and said second proxy residing on said second computer, said first proxy containing network linkage and indication to access said programmed methods on said second computer and said second proxy containing network linkage and indication to access said programmed methods on said first computer; and,

computer-readable program code means for accessing said remote programmed methods through said proxies.

5. A computer readable storage medium according to Claim 4 wherein said second proxy contains similarly named methods to the those methods in the objects remaining on said first computer and wherein said first proxy contains similarly named methods to those methods in the objects moved to said second computer.

6. A computer readable storage medium according to Claim 4 or 5 wherein said network linkage includes statements for remote invocation of the programmed methods.

7. A system for distributing objects of a program, each object containing one or more programmed methods, across more than one physical device, said system comprising :

means for identifying all of the objects in the program;

means for determining which of the objects are to reside on a first computer and which of the objects are to reside on a second computer;

means for identifying all programmed methods contained in each object that may be accessed from a remote computer;

means for generating a first proxy and a second proxy for each object that may be accessed from a remote computer, said first proxy residing on said first computer and said second proxy residing on said second computer, said first proxy containing network linkage and indication to access said programmed methods on said second computer and said second proxy containing network linkage and indication to access said programmed methods on said first computer; and,

means for accessing said remote programmed methods through said proxies.

8. A system as claimed in Claim 7 wherein said second proxy contains similarly named methods to the those methods in the objects remaining on

said first computer and wherein said first proxy contains similarly named methods to those methods in the objects moved to said second computer.

9. A system as claimed in Claim 7 or 8 wherein said network linkage includes statements for remote invocation of the programmed methods.

5



Application No: GB 9808749.7
Claims searched: 1-3 & 7-9

Examiner: Melanie Gee
Date of search: 7 October 1998

Patents Act 1977 Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK Cl (Ed.P): G4A (APX)

Int Cl (Ed.6): G06F 9/46

Other: Online: WPI, INSPEC, COMPUTER

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
A	GB 2305270 A (INTERNATIONAL BUSINESS MACHINES), see pages 1-6.	
A	US 5684955 A (MEYER et al.), see whole document.	

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

THIS PAGE BLANK (USPTO)